



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

mv

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/688,070	10/17/2003	Douglas P. Doucette	103.1045.04	8214
22883	7590	11/08/2006	EXAMINER	
SWERNOFSKY LAW GROUP PC P.O. BOX 390013 MOUNTAIN VIEW, CA 94039-0013			MOFIZ, APU M	
			ART UNIT	PAPER NUMBER
			2165	

DATE MAILED: 11/08/2006

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary

Application No.

10/688,070

Applicant(s)

DOUCETTE ET AL.

Examiner

Apu M. Mofiz

Art Unit

2165

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 29 September 2006.
- 2a) ☒ This action is **FINAL**. 2b) ☐ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-7, 10, 12, 14-18 and 27-30 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-7, 10, 12, 14-18 and 27-30 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☐ The drawing(s) filed on _____ is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.
- Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
- Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
 2. ☐ Certified copies of the priority documents have been received in Application No. _____.
 3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- | | |
|--|---|
| 1) <input type="checkbox"/> Notice of References Cited (PTO-892) | 4) <input type="checkbox"/> Interview Summary (PTO-413) |
| 2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948) | Paper No(s)/Mail Date. _____ |
| 3) <input type="checkbox"/> Information Disclosure Statement(s) (PTO/SB/08) | 5) <input type="checkbox"/> Notice of Informal Patent Application |
| Paper No(s)/Mail Date _____ | 6) <input type="checkbox"/> Other: _____ |

DETAILED ACTION

Response to Arguments

1. Applicant's arguments submitted on 09/29/2006 with respect to claims 1-7, 10, 12, 14-18 and 27-30 have been reconsidered but are not deemed persuasive for the reasons set forth below.

Examiner's Responses to Applicant's Remarks are listed below:

2. Applicant argues (under REMARKS section) that, Hitz does not disclose:

(a) A method for improved space allocation in a file system having a set of storage blocks in a mass storage system, including maintaining an active map of said storage blocks not available for writing data (i.e., **"To implement consistency points, WAFL always writes new data to unallocated blocks on disk.** (col 4, lines 13-16)

... The present invention also creates **snapshots, which are virtual read-only copies of the file system.** (col 4, lines 20-30) ... The present invention prevents new data written to the active file system from overwriting "old" data that is part of a snapshot(s). It is necessary that old data not be overwritten as long as it is part of a snapshot. **This is accomplished by using a multi-bit free-block map.** ... A first bit indicates whether a block is used by the active file system, and 20 remaining **bits are used for up to 20 snapshots**, however, some bits of the 31 bits may be used for other purposes. (col 4, lines 32-45) ... The present invention uses a Write Anywhere File-system Layout (WAFL). **This disk format system is block based (i.e., 4 KB blocks that have no fragments)**, uses inodes to describe its files, (col 5, lines 48-52)

... Another meta-data file is the “**block map**” (blkmap) file. FIG. 11A is a diagram illustrating a blkmap file 1110. The blkmap file 1110 contains a 32-bit entry 1110A-1110D for each 4 KB block in the disk system. **It also serves as a free-block map file.** (col 6, lines 50-55) ... the entry 1110A of blkmap file 1110 indicates a block that is part of the **active file system**. Bits 1-20 (BIT1-BIT20) are used to indicate corresponding snapshots, if any, that reference the block. (col 10, lines 8-11)” The preceding text excerpts and Fig. 11 clearly indicate that the Write Anywhere File System (WAFL) tries to maintain consistency in the file system. It divides the disk into regions of 4 KB blocks to write data (can be user data or any other data; any file system including WAFL always write data i.e., user data to disk). But to maintain consistency in writing to the disk, it maintains a blkmap file (an active map or a meta data file containing binary data; the system maintains also some other metadata files), which contains the listings for available blocks (blocks are of equal sizes of 4 KB) and snapshot of the file system (i.e., a copy of the file system at a previous time). The actual available disk blocks are all storage blocks that stores any data that any file system allows to write on a disk (e.g. a user file). Therefore if the file system needs to write a file size of 64 KB, it queries/checks the blkmap file (and/or other metadata files) to determine what blocks are available to be written. Therefore the values i.e., the available/free block information (i.e., the active map file/ blkmap file contains entries for each 4KB block in the disk system and also indicates if the blocks are free and also snapshots of the file system) is queried/ checked to determine if/where the file can be written. The blkmap (i.e., the active map of the active file system) itself may be meta data, but certainly the blocks that the blkmap file keeps track of are not for meta data only (i.e., these are just plain storage blocks like any other file system and can store any data e.g., user data, meta data, picture data or any other

Art Unit: 2165

data that the computer allows). The blkmap file contains information/value for the storage blocks, but the blkmap file itself is not for writing/storing user data.), including maintaining an active map of said storage blocks not available for writing data (See citations and explanations above); determining, for each one of a plurality of equal regions of said storage blocks in said mass storage system, a value indicative of a number of storage blocks available for writing data in said region, based on said active map and at least one snapshot of the file system, each said at least one snapshot of the file system having a copy of said active map at a previous time (See citations and explanations above); and selecting, based on the values, at least one of said plurality of regions for writing data (See citations and explanations above).

Any other arguments by the applicant are more limiting or irrelevant than the claimed language.

Claim Rejections - 35 USC § 102

3. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

4. Claims 1,6,7,14,16 and 27-30 are rejected under 35 U.S.C. 102(b) as being anticipated by Hitz et al. (U.S. Patent No. 5,819,292 and Hitz hereinafter).

As to claims 1,14,16 and 27-30, Hitz teaches a method for improved space allocation in a file system having a set of storage blocks in a mass storage system, including maintaining an active map of said storage blocks not available for writing data (i.e., **“To implement consistency points, WAFL always writes new data to unallocated blocks on disk.** (col 4, lines 13-16) ... The present invention also creates **snapshots, which are virtual read-only copies of the file system.** (col 4, lines 20-30) ... The present invention prevents new data written to the active file system from overwriting “old” data that is part of a snapshot(s). It is necessary that old data not be overwritten as long as it is part of a snapshot. **This is accomplished by using a multi-bit free-block map.** ... A first bit indicates whether a block is used by the active file system, and 20 remaining **bits are used for up to 20 snapshots**, however, some bits of the 31 bits may be used for other purposes. (col 4, lines 32-45) ... The present invention uses a Write Anywhere File-system Layout (WAFL). **This disk format system is block based (i.e., 4 KB blocks that have no fragments)**, uses inodes to describe its files, (col 5, lines 48-52) ... Another meta-data file is the **“block map”** (blkmap) file. FIG. 11A is a diagram illustrating a blkmap file 1110. The blkmap file 1110 contains a 32-bit entry 1110A-1110D for each 4 KB block in the disk system. **It also serves as a free-block map file.** (col 6, lines 50-55) ... the entry 1110A of blkmap file 1110 indicates a block that is part of the **active file system**. Bits 1-20 (BIT1-BIT20) are used to indicate corresponding snapshots, if any, that reference the block. (col 10, lines 8-11)” The preceding text excerpts and Fig. 11 clearly indicate that the Write

Art Unit: 2165

Anywhere File System (WAFL) tries to maintain consistency in the file system. It divides the disk into regions of 4 KB blocks to write data (can be user data or any other data; any file system including WAFL always write data i.e., user data to disk). But to maintain consistency in writing to the disk, it maintains a blkmap file (an active map or a meta data file containing binary data; the system maintains also some other metadata files), which contains the listings for available blocks (blocks are of equal sizes of 4 KB) and snapshot of the file system (i.e., a copy of the file system at a previous time). The actual available disk blocks are all storage blocks that store any data that any file system allows to write on a disk (e.g. a user file). Therefore if the file system needs to write a file size of 64 KB, it queries/checks the blkmap file (and/or other metadata files) to determine what blocks are available to be written. Therefore the values i.e., the available/free block information (i.e., the active map file/ blkmap file contains entries for each 4KB block in the disk system and also indicates if the blocks are free and also snapshots of the file system) is queried/ checked to determine if/where the file can be written. The blkmap (i.e., the active map of the active file system) itself may be meta data, but certainly the blocks that the blkmap file keeps track of are not for meta data only (i.e., these are just plain storage blocks like any other file system and can store any data e.g., user data, meta data, picture data or any other data that the computer allows). The blkmap file contains information/value for the storage blocks, but the blkmap file itself is not for writing/storing user data. The purpose of the Write Anywhere File System is to write data (i.e., any data including user data) to disk storage blocks consistently, and the metadata files e.g., blkmap file is to help achieve (determine) its goal.), including maintaining an active map of said storage blocks not available for writing data (See citations and explanations above); determining, for each one of a plurality of equal regions

Art Unit: 2165

of said storage blocks in said mass storage system, a value indicative of a number of storage blocks available for writing data in said region, based on said active map and at least one snapshot of the file system, each said at least one snapshot of the file system having a copy of said active map at a previous time (See citations and explanations above); and selecting, based on the values, at least one of said plurality of regions for writing data (See citations and explanations above).

As to claim 6, Hitz teaches that wherein each said value is a binary number (See citations and explanations in claim 1 rejection above) (Fig. 11A-D; Abstract; col 4, lines 5-45; col 9, lines 50-67; col 10, lines 1-19; col 11, lines 1-67; col 12, lines 1-67).

As to claim 7, Hitz teaches that wherein each said value stored in a data block containing one or more binary numbers each corresponding to a unique region (See citations and explanations in claim 1 rejection above) (Fig. 11A-D; Abstract; col 4, lines 5-45; col 9, lines 50-67; col 10, lines 1-19; col 11, lines 1-67; col 12, lines 1-67).

Claim Rejections - 35 USC § 103

5. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

6. Claims 2-5,10,12,15 and 17-18 are rejected under 35 U.S.C. 103(a) as being unpatentable over Hitz et al. (U.S. Patent No. 5,819,292 and Hitz hereinafter) in view of Hitz et al. (U.S. Patent No. 6,038,570 and Hitz_2 hereinafter).

As to claims 2 and 15, Hitz teaches selecting based on the values at least one of plurality of regions for writing data (See citations and explanations in claim 1 rejection above) (Fig. 11A-D; Abstract; col 4, lines 5-45; col 9, lines 50-67; col 10, lines 1-19; col 11, lines 1-67; col 12, lines 1-67).

Hitz does not teach setting an allocation threshold and comparing the values to the threshold.

Hitz_2 teaches setting an allocation threshold and comparing the values to the threshold (col 20, lines 19-30).

It would have been obvious to a person of ordinary skill in the art at the time of Applicant's invention to modify the teachings of Hitz with the teachings of Hitz_2 to include setting an allocation threshold and comparing the values to the threshold with the motivation to be very efficient to write to stripes in the RAID array where there are no allocated blocks in a stripe on the data disks (Hitz_2, col 20, lines 27-30).

As to claims 3 and 16, HITZ teaches writing the data into the selected at least one of said plurality of regions (See citations and explanations in claim 1 rejection above) (Fig. 11A-D; Abstract; col 4, lines 5-45; col 9, lines 50-67; col 10, lines 1-19; col 11, lines 1-67; col 12, lines 1-67).

As to claims 4 and 17, HITZ_2 teaches setting the threshold based on percentage of the number of storage blocks available for writing data in the file system (col 20, lines 19-30).

It would have been obvious to a person of ordinary skill in the art at the time of Applicant's invention to modify the teachings of Hitz with the teachings of Hitz_2 to include setting the threshold based on percentage of the number of storage blocks available for writing data in the file system with the motivation to be very efficient to write to stripes in the RAID array where there are no allocated blocks in a stripe on the data disks (Hitz_2, col 20, lines 27-30).

As to claims 5 and 18, Hitz teaches selecting based on the values at least one of plurality of regions for writing data (See citations and explanations in claim 1 rejection above) (Fig. 11A-D; Abstract; col 4, lines 5-45; col 9, lines 50-67; col 10, lines 1-19; col 11, lines 1-67; col 12, lines 1-67).

Hitz does not teach using threshold for selecting a region.

Hitz_2 teaches using threshold for selecting a region (col 20, lines 19-30).

It would have been obvious to a person of ordinary skill in the art at the time of Applicant's invention to modify the teachings of Hitz with the teachings of Hitz_2 to include using threshold for selecting a region with the motivation to be very efficient to write to stripes in the RAID array where there are no allocated blocks in a stripe on the data disks (Hitz_2, col 20, lines 27-30).

As to claim 10, Hitz teaches said selecting comprises linearly searching said plurality of regions to select a first region based on corresponding values (See citations and explanations in claim 1 rejection above) (Fig. 11A-D; Abstract; col 4, lines 5-45; col 9, lines 50-67; col 10, lines 1-19; col 11, lines 1-67; col 12, lines 1-67).

Hitz does not teach using threshold for selecting a region.

Hitz_2 teaches using threshold for selecting a region (col 20, lines 19-30).

It would have been obvious to a person of ordinary skill in the art at the time of Applicant's invention to modify the teachings of Hitz with the teachings of Hitz_2 to include using threshold for selecting a region with the motivation to be very efficient to write to stripes in the RAID array where there are no allocated blocks in a stripe on the data disks (Hitz_2, col 20, lines 27-30).

As to claim 12, Hitz teaches further including additional selecting when said data requires more blocks than available in the selected at least one of said plurality of equal regions (See citations and explanations in claim 1 rejection above) (Fig. 11A-D; Abstract; col 4, lines 5-45; col 9, lines 50-67; col 10, lines 1-19; col 11, lines 1-67; col 12, lines 1-67).

Conclusion

7. **THIS ACTION IS MADE FINAL.** See MPEP § 706.07(a). Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).


A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the date of this final action.

Points of Contact

6. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Apu M. Mofiz whose telephone number is (571) 272-4080. The examiner can normally be reached on Monday – Thursday 8:00 A.M. to 4:30 P.M.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Jeffrey Gaffin can be reached at (571) 272-4146. The fax numbers for the group is (571) 273-8300.

Any inquiry of a general nature or relating to the status of this application should be directed to the Group receptionist whose telephone number is (703) 305-9600.



Apu M. Mofiz
Primary Patent Examiner
Technology Center 2100

November 06, 2006